

Attorney Ref.: 42P9347C
Express Mail No.: EV339914529US

UNITED STATES PATENT APPLICATION

FOR

**METHOD AND APPARATUS FOR A LOW
LATENCY SOURCE-SYNCHRONOUS ADDRESS
RECEIVER FOR A HOST SYSTEM BUS IN A MEMORY
CONTROLLER**

Inventors:

**Srinivasan T. Rajappa
Romesh B. Trivedi
Rajagopal Subramanian
Zohar Bogin
Serafin Garcia**

Prepared By:

**BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN, LLP
12400 Wilshire Blvd., 7th Floor
Los Angeles, California 90025-1030
(310) 207-3800**

**METHOD AND APPARATUS FOR A LOW
LATENCY SOURCE-SYNCHRONOUS ADDRESS
RECEIVER FOR A HOST SYSTEM BUS IN A MEMORY
CONTROLLER**

RELATED APPLICATIONS

[0001] This utility application is a continuation of U.S. Application No. 09/665,922, filed on September 20, 2000, currently pending.

FIELD OF THE INVENTION

[0002] The present invention relates generally to memory controllers. In particular, the present invention relates to an apparatus and method for a low latency source-synchronous address receiver for a system bus in a memory controller.

BACKGROUND OF THE INVENTION

[0003] Many computer devices operate based on an external clock. For example, a processor may receive a clock input and perform all operations or events only when the clock transitions. Devices in which events proceed based on a clock transition are referred to as "synchronous" devices.

[0004] Other computer devices do not base their operation on an external clock. These devices are referred to as "asynchronous" or "self-timed" devices. A self-timed device typically receives a request from a processor. The device then performs the operation and indicates to the processor when the operation is complete. However, the time required for the operation to complete is not based on an external clock (i.e., a predetermined number of clock cycles). Rather, in the case of a self-timed device, the time required is based on the asynchronous delay paths through the device, which may vary in duration based on the operations that are performed.

[0005] In a conventional memory controller architecture, the memory controller architecture generally includes a processor, a chipset and a main memory. A host system bus which connects the processor to the chipset is a synchronous device generally controlled by a common clock interface. In other words, the speed at which the host system bus can run is limited by the speed of the system clock. As technology pushes the processing speed, common clock interface buses run the risk of creating a bottleneck in memory controller architectures. In fact, these advances in processor design have pushed memory controller systems to a level where the speed of a bus or an architecture cannot be scaled using an increased clock frequency. One technique for accommodating the increased processor speed is to replace the host system bus with a source synchronous system bus.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006] The features, aspects, and advantages of the embodiments described herein will become more fully apparent from the following detailed description and appended claims when taken in conjunction with accompanying drawings in which:

[0007] FIG. 1 depicts a block diagram of a memory controller architecture system, in accordance with one embodiment;

[0008] FIG. 2 depicts a block diagram further illustrating the address receiver of FIG. 1, in accordance with one embodiment;

[0009] FIG. 3 depicts a timing diagram illustrating the functionality of the receiver as depicted in FIG. 2, in accordance with one embodiment;

[00010] FIG. 4 depicts a diagram of a source synchronous address receiver according to one embodiment;

[00011] FIG. 5 depicts a block diagram of a conventional memory controller architecture system according to one embodiment;

[00012] FIG. 6 depicts the source synchronous address receiver according to embodiment; and

[00013] FIG. 7 depicts a diagram illustrating the functionality of the source synchronous address receiver as depicted in FIG. 6, in accordance with one embodiment.

DETAILED DESCRIPTION OF THE INVENTION

[00014] An apparatus and method for a low latency source synchronous address receiver for a system bus in a memory controller are described. In one embodiment, a low latency path between a system bus address input to a memory bus is provided, resulting in a high performance memory controller by using a flow-through path. In one embodiment, flow-through path is controlled by two inputs: one is a source synchronous strobe directing the address receiver to latch an address and store it, while the other is a protocol signal that signals the beginning of an address transfer which enables the flow-through path.

[00015] In the following detailed description, numerous specific details are set forth to provide a thorough understanding of the embodiments described herein. However, one having ordinary skill in the art should recognize that the invention may be practiced without these specific details. For example, various signals, layout patterns, memory cell configurations and circuits, and logic circuits may be modified according to the embodiments described herein. The following description provides examples, and the accompanying drawings show various examples for the purposes of illustration. However, these examples should not be construed in a limiting sense as they are merely intended to provide examples of the at least one embodiment, rather than to provide an exhaustive list of all possible implementations of the present invention. In some instances, well-known structures, devices, and techniques have not been shown in detail to avoid obscuring the present invention.

[00016] The following system architecture describes specific embodiments for implementing a double-pumped, source synchronous address receiver for a system bus. However, those skilled in the art will appreciate that the embodiments may be implemented using various circuit design modifications. Specifically, the flow-through path as taught by the present invention can be implemented using various logic design techniques while remaining within the scope of the embodiments described herein. Moreover, although one embodiment describes a two input, source synchronous address receiver, those skilled in the art will realize that the embodiments described herein can be easily extended to higher order interfaces by scaling the receiver structure.

[00017] In a conventional memory controller architecture, the memory controller generally includes a host system bus, which connects the processor to a chipset. Generally, the host system is controlled by a common clock interface. In other words, the speed at which the host bus can run is limited by the speed of the system clock. As technology pushes the processing speed of CPUs, common clock interface buses run the risk of creating a bottleneck in memory controller architectures. In fact, these advances in processor design have pushed memory controller systems to a level where the speed of a bus or an architecture cannot be scaled using an increased clock frequency.

[00018] Accordingly, in one embodiment, a memory controller architecture system 100 is provided using a source synchronous system bus 110. Representatively, system

100 comprises source synchronous processor system bus (front side bus (FSB)) 110 for communicating information between processor (CPU) 102 and chipset 104. As described herein, the term "chipset" is used in a manner to collectively describe the various devices coupled to CPU 102 to perform desired system functionality.

[00019] In one embodiment, chipset 104 includes a memory controller to communicate with main memory 106. In one embodiment, main memory 106 may include, but is not limited to, random access memory (RAM), dynamic RAM (DRAM), static RAM (SRAM), synchronous DRAM (SDRAM), double data rate (DDR) SDRAM (DDR-SDRAM), Rambus DRAM (RDRAM) or any device capable of supporting high-speed buffering of data.

[00020] Representatively, chipset 104 may include an address receiver for receiving an address packet from CPU 110 in order to signal an address transaction request. Once the address packet is received by chipset 104, the chipset decodes the address packet 122 to generate an address of the requested data in main memory and return the requested data to the CPU. In one embodiment, address receiver 120 is further illustrated with reference to FIG. 2.

[00021] Representatively, address packet 122 may include a phase zero component (IFA[0]) 132, as well as phase one information (IFA[1]) 140. Unfortunately, phase zero information 132 of address packet 122 is vital for decoding of address packet 122 by chipset 104. Accordingly, a memory controller system, such as memory controller system 100 as depicted in FIG. 1, requires a low latency path between an address input of FSB 110 and memory bus 108 in order to achieve high performance unless phase zero address 132 can be instantly provided to core chipset logic 104. Once the address packet appears on an address pin, any benefits provided by using source synchronous FSB 110 are lost.

[00022] As depicted in FIG. 3, the address packet 122 includes a first component or phase zero component 132 and a second component or phase one component 140. The phase zero component 132 of the address packet 122 describes a transaction type of the address packet 122 (transaction address), indicating whether the address of the requested data is memory based or IO based. Consequently, this information is crucial to the memory controller architecture 100 for high-speed decoding of the address packet 122 by the chipset 104. The phase one component 140 of the address packet 122 contains address attributes of the address packet including data size/length attributes, byte enables, defer ID, extended functions, cycle types, etc. The information provided by the phase one component 140 of the address packet 122 is not as critical to the decoding stages as the phase zero component 114 of the address packet 112.

[00023] Referring again to FIG. 2, phase zero component 132 of the address packet 122 appears at an output 130 of first flip-flop 126 in response to the falling edge of strobe signal 124. Likewise, phase one component 140 of address packet 122 appears at an output 138 of the second flip-flop 134 in response to the rising edge of the strobe signal

124. Unfortunately, the address receiver 120 as depicted in FIG. 3 introduces a delay (TDA) 156, as indicated in the timing diagram, of as much as a few nanoseconds in producing the phase zero component 132 of the address packet 122 at the output 130 of the first flip-flop 126.

[00024] Accordingly, FIG. 4 depicts an embodiment illustrating a source synchronous address receiver (SSAR) 200 for use within a chipset memory controller system architecture, such as in a host system bus of a memory controller architecture 300 as depicted in FIG. 5. The source synchronous address receiver 200 includes an input differential amplifier 202 that compares an address packet 122 received on an address pin 204 for a data request against a reference voltage 206 to generate a digital address packet as the address packet 122. The address packet 122 includes a first, or phase zero component 132 (IFA[0]), describing a transaction type of the address packet 122 (transaction address) and a second, or phase one component 140 (IFA[1]), describing attributes of the address packet 122 as described above with reference to FIG. 3.

[00025] In one embodiment, flow-through circuit 220 receives a digital address strobe signal 210 and a digital address select signal 208. The digital address strobe signal 210 is a digital version of an analog source synchronous strobe signal 124 that directs the address receiver 200 to latch and store address information available on an address system bus such as the system bus 310. The digital address select signal 208 is generated from an analog common clock protocol signal 154 that signals the beginning of an address transfer and is used to enable a flow-through path as described below.

[00026] In one embodiment, flow-through circuit 220 generates an enable signal 240 in response to the digital address strobe signal 210 and the digital address select signal 208. The enable signal 240 is then provided to a flow-through gate 242 having the address packet as an input 244. The flow-through gate 242 provides the first or phase zero component 132 of the address packet 122 (transaction address) to a chipset, such as chipset 304, once the address packet 122 appears on the address pin 204. In one embodiment, the flow-through gate 242 provides a flow-through path from the address pin 204 to the chipset 304 for the transaction address 122 to expedite the initiation of decoding of the address packet 122 by the chipset 304. Representatively, first flip-flop 250 receives the digital address packet 122 and the digital address strobe signal 210 as inputs and provides the second or phase one component 140 of the address packet 122 to the chipset in response to the address strobe signal 208. Once the second, or phase one component 140, of the address packet 122 is provided to the chipset 304, the chipset 304 can complete decoding of the address packet 122.

[00027] FIG. 6 is a block diagram depicting the source synchronous address receiver 200 according to one embodiment. The flow-through circuit 220 further includes a second flip-flop 222, including a feedback inverter 224 coupled between an input 226 and an output 228 of the second flip-flop 222. The second flip-flop 222 also includes the

digital address select signal 208 as a clock pulse input. A third flip-flop 230 also includes a feedback inverter 232 coupled between an input 234 and an output 236 of the third flip-flop 230. The third flip-flop 230 also includes the digital address strobe signal 210 as a clock pulse input. An exclusive-OR gate 238 includes the output of the second flip-flop 222 and the output of the third flip-flop 230 as inputs to generate the enable signal 240 for the flow-through gate 242.

[00028] In one embodiment, address receiver 200 also includes a first differential amplifier 260 that compares the analog common clock protocol signal 154, received on an address select pin 262, against the reference voltage 206 to generate a digital common clock protocol signal 264. A fourth flip-flop 266 receives the digital common clock protocol signal into signal 264 as an input 268 and a common clock signal (ABUTFCLK 100) as a clock pulse input 270. Once received, the fourth flip-flop 266 generates a flopped address select signal 216 at an output 272. An inverter 274 then receives the flopped address select signal 216 and generates the digital address select signal 208 at an output 276 of the inverter 274 for input to the flow-through circuit 220.

[00029] The source synchronous address receiver 200 also includes a second differential amplifier 280 that compares an analog address strobe signal 124 received on a address strobe pin 282 against the reference voltage 206 to generate the digital address strobe signal 210. In one embodiment, source-synchronous address receiver 200 is used for the system address bus 310 of the memory controller 300 as depicted in FIG. 5. In one embodiment, CPU 302 of the memory controller 300 is preferably a Willamette® generation CPU as manufactured by the Intel Corporation. The flow-through gate 242 is preferably a latch, although various other logic gates are within the contemplation of the described embodiments. In addition, the first flip-flop 250, the second flip-flop 222, the third flip-flop 230, and the fourth flip-flop 266 are preferably data flip-flops, although various types of logic gates are within the contemplation of the described embodiments..

[00030] One embodiment illustrating operation of the source synchronous address receiver 200 is described with reference to FIGS. 6 and 7. Referring to FIG. 7, a timing diagram 400 is depicted illustrating one embodiment. An address packet 122 (pad) as depicted in FIG. 6 illustrates the first component, or phase zero component 132 (A0), and the second, or phase one component 140 (A1). The enable signal 240 (POLENB) is initially low. Since the flow-through gate 242 is active low enabled, during reset of the enable signal 240, the flow-through gate is enabled or active low, thus making the flow-through gate 242 transparent.

[00031] Accordingly, in one embodiment, receiver 200 enables the flow-through path from the address pin 204 to the core logic of the chipset 304 (FIG. 45) for the first component, or phase zero component 132, of the address packet 122. In one embodiment, at the moment the phase zero component 132 of the address packet 122 appears on the address pin 204, the data is immediately made available to the chipset as IFA[0]. Enabling

of the flow-through gate 242 is toggled by the falling edge of the digital source synchronous address strobe signal 210 (ABUTFADSTB), thus closing the flow-through gate 242 as indicated by the arrow 402. The flow-through gate 242 is again made transparent by the falling edge of the digital address select signal 208 (ABUTADS), thus enabling the flow-through path for next phase zero address component 132 of the address packet 122 (B0) as indicated by the arrow 404.

[00032] The phase one component 140 of the address packet 122 (A1) is sampled at the rising edge of the digital address strobe signal 210 and appears as IFA[1] as indicated by the arrow 406. The digital address select signal 208 (ABUTFADS) is essentially a complement of the flopped address select signal 216 (flopped ADS#) as described with reference to FIG. 6 which is generated from the digital common clock protocol signal 264. Once the phase one component 140 of the address packet 122 is received by the chipset 304 (FIG. 5), the chipset 304 (FIG. 5) can complete decoding of the address packet 122.

[00033] In one embodiment, once the address packet is decoded by the chipset 304, the decoded address is provided to the main memory 306 via the memory bus 308 to retrieve the data requested by the CPU 302 and consequently transferred to the CPU 302 as depicted in FIG. 5. Consequently, in one embodiment, a flow-through path achieves a low-latency path between the system bus address input to the memory bus resulting in a high performance memory controller, which can accommodate the increased CPU processing speeds required in today's technology, as well as future processing speeds.

[00034] In one embodiment, a source synchronous address receiver provides a low latency path between the host system bus and the memory bus to enable a high performance memory controller. In one embodiment, this low latency path is achieved by using a flow-through path. This flow-through path is controlled by two inputs, one is a source synchronous strobe directing the apparatus to latch an address and store it, while the other is a common clock protocol signal that signals the beginning of the address transfer which enables the flow-through path. This use of source synchronous and common clock signals to achieve the desired end result (minimal latency to the memory address) is unique to the described embodiment. The embodiments described herein may be used in future implementations of source synchronous system address buses and memory controller devices.

[00035] Having disclosed exemplary embodiments and the best mode, modifications and variations may be made to the disclosed embodiments while remaining within the scope of the invention as defined by the following claims.